


EXHIBIT 11

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx
Go
SEP
NOV
JAN
29
125 captures
12 Feb 2007 - 8 Feb 2021
2013
2014
2015
About this capture



- Home
- Learning
 - Tutorials
 - Data Structures
 - Andersson Trees
 - AVL Trees
 - Binary Search Trees I
 - Binary Search Trees II
 - Hash Tables
 - Red Black Trees
 - Skip Lists
 - Linked Lists
 - Algorithms
 - Languages
 - Articles
 - Asymptotic Notation
 - Using rand()
 - Libraries
- Licensing

Welcome back! Or if this is your first experience with my tutorials, get ready for a good time. But first, why another red black tree tutorial? Anyone who searches for red black trees on Google will be rewarded with a slew of resources that include tutorials, general descriptions, Java applets, papers, libraries, and power point presentations. This is good, but not good enough. Most resources are content with only looking at insertion (slackers!). Others fail to describe deletion in such a way that one can easily understand why it works, obviously thinking that enumerating the cases should be enough. Even worse, every resource fails to notice that red black trees can be written several different ways.

Well guys, I'm not that smart. Insertion into a red black tree is painfully simple. Any programmer with half a brain and a little experience with binary search trees can figure it out with minimal effort. Deletion on the other hand is, quite frankly, a pain in the ass. I had better luck independently rediscovering the cases than trying to figure out why some of the traditional cases are used. For example, a sibling should be black, so a special case makes it so, yet nobody seems to know why. After dealing with that case, I can both tell you and show you why (it's not called a case reduction for nothing).

When it comes to red black trees, there is only one widely accepted resource: "Introduction to Algorithms" by Cormen, Leiserson, and Rivest, affectionately known as CLR. This is a good resource because not only does it describe both insertion and deletion, it gives sufficient pseudocode for both operations! As the only well known and easily accessible text that does this, just about every red black tree in the real world uses a translation of the CLR algorithms. Unfortunately, CLR's approach is both complicated and inefficient, through heavy use of parent pointers. Red black trees don't have to be that hard.

I learned red black trees the hard way. Given only a description of the rules, I worked through every case I could think of while trying to devise an algorithm. Without a copy of CLR (I still don't own it), I wasn't able to see how the problem was meant to be solved. This is a good thing because I "discovered" several ways of doing it, and I believe that I peeled away most of the complexity. I cut myself on the sharp edges of the algorithms, I wept from frustration trying to get them to work, and now I'll share the results of that effort so that you don't have to go through the same ordeal.

This tutorial will cover both the non-recursive top-down and recursive bottom-up algorithms for red black insertion and deletion, as well as provide full C source code for all variations described. The most common implementation is based off of the pseudocode from CLR, but I won't cover that because it's just another implementation of the bottom-up process. If you're interested in how to implement red black trees with parent pointers, either CLR or Ben Pfaff's excellent online book will serve you well. If you're interested in a non-recursive red black implementation using an explicit stack instead of recursion, Ben Pfaff's online book also covers this. I see no point in reproducing something that is dealt with adequately elsewhere.

http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx
Go
SEP
NOV
JAN
29
125 captures
12 Feb 2007 - 8 Feb 2021
2013
2014
2015
About this capture

Basic binary search trees are simple data structures that boast $O(\log N)$ search, insertion, and deletion. For the most part this is true, assuming that the data arrives in random order, but basic binary search trees have three very nasty degenerate cases where the structure stops being logarithmic and becomes a glorified linked list. The two most common of these degenerate cases is ascending or descending sorted order (the third is outside-in alternating order). Because binary search trees store their data in such a way that can be considered sorted, if the data arrives already sorted, this causes problems. Consider adding the values 0,1,2,3,4 to a binary search tree. Since each new item is greater in value than the last, it will be linked to the right subtree of every item before it:

```


c
 \
  1
   \
    \

```


Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10



That's not a very good binary search tree. The performance degrades from $O(\log N)$ to $O(N)$ because the tree is now effectively a linear data structure. We would rather have two choices at each node instead of just one. This way we can take full advantage of the two-way design of binary search trees:



Many brain cells have perished trying to come up with an easy way to guarantee this optimum structure, which is called a balanced binary search tree. Unfortunately, it takes far too much work to maintain a perfectly balanced binary search tree, so it's impossible to efficiently guarantee perfect balance at all times. However, we can come close enough to guarantee logarithmic performance. Several of these "close enough" balancing schemes are in common use, and the two forerunners are AVL trees and red black trees. Because I have another tutorial on AVL trees, we will only cover red black trees in detail here.

Red black trees were originally an abstraction of an abstraction that managed to become a concrete data structure of their own. The original abstraction was suggested by Rudolf Bayer, which he called symmetric binary B-trees. The idea was to simulate a B-tree of order k by using k binary trees.

http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

125 captures

12 Feb 2007 - 8 Feb 2021

Go

SEP

SEP

JAN

29

not a binary

tree,

2013

2021

2015



The basic idea behind the symmetric binary B-tree is that a node can have horizontal or vertical links. In this way a binary search tree can simulate the structure of B-tree nodes. A vertical link separates two different nodes and a horizontal link separates nodes that are treated logically as the same B-tree node. The equivalent B-tree and symmetric binary B-tree (SBB-tree) nodes are as follows (* = Unknown node). Notice how the B-tree node's structure is faked using multiple binary tree nodes:


2-node

SBB-tree B-tree




3-node

SBB-tree 1 SBB-tree 2 B-tree



4-node


SBB-tree 1 B-tree



Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10



The algorithms for maintaining B-trees are confusing at best, probably because they grow up instead of down and programmers have trouble with that. As a result, the algorithms for B-trees don't translate well into the symmetric binary B-tree abstraction and are still confusing. So while it was a good idea, symmetric binary B-trees were missing something.

Later, Robert Sedgwick and Leonidas Guibas came up with a mnemonic abstraction that made symmetric binary B-trees easier to understand (and the name was shorter too!). By giving nodes a color, you could easily differentiate between vertical and horizontal links. Under this abstraction, nodes that are part of a logical B-tree node (horizontal links) are colored red, while nodes that separate logical B-tree nodes (vertical links) are colored black. Thus was born the red black tree. Here are the 2, 3, and 4-nodes

http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtrees.aspx

125 captures

12 Feb 2007 - 8 Feb 2021

Go

SEP

SEP

JAN


29

2013

2021

2015

ABOUT THIS CAPTURE



This gives us some easy to remember rules to maintain balance. First, a node that's colored red can't have a child that's also colored red, because that doesn't fit into the abstraction. If a node is colored red, then it's part of a logical node, and the only possible color for the next node is black because a red node's links play the part of a B-tree node's links to the next B-tree node. A violation of this rule is called a red violation (duh!).

Next, because in a B-tree any path from the root to a leaf always has the same number of nodes, and all leaves are on the same level, the number of black nodes along any path in a red black tree must be the same. Because a black parent and its red children are a part of the same logical node, red nodes aren't counted along the path. This is called the black height of a red black tree. A violation of this rule is called a black violation.

Finally, the root of a red black tree must be black. This makes all kinds of sense because only a black node can have red children, and the root has no parent. However, this rule only applies when using the B-tree abstraction. In practice the root can be either red or black and no affect the performance of a red black tree.

All algorithms that maintain balance in a red black tree must not violate any of these rules. When all of the rules apply, the tree is a valid red black tree, and the height cannot be any shorter than $\log(N + 1)$ but also no taller than $2 * \log(N + 1)$. This is obviously logarithmic, so a red black tree guarantees approximately the best case for a binary search tree.

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsv_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

Now, because red black trees are so common these days, the original abstraction has been largely forgotten in favor of a data structure that simply avoids violating two basic

of black nodes. This fulfills the requirements of a red black tree n

does make the algorithm simpler because you can avoid the special case of a red violation at the root.

Before we start, we'll need to cover some of the basics and get a little framework code going. First, the node and tree structures. However you want to represent a node's color is up to you. A common solution is an enumeration with with two constants, one for black and one for red. I find it easier to use a flag where if the flag is set, the node is red and if the flag is not set, the node is black. Naturally the inverse works as well, but the flag solution with red being set will be what this tutorial uses:

```

1 struct jsw_node {
2     int red;
3     int data;
4     struct jsw_node *link[2];
5 };
6
7 struct jsw_tree {
8     struct jsw_node *root;
9 };

```

I'm sure that if this is your first time reading one of my tree tutorials that the link array could be confusing. Instead of using two pointers called left and right, as is the usual convention, I prefer to use an array where index 0 is left and index 1 is right. This makes it easier to merge the symmetric cases together to shorten the code and to guarantee correctness. This is my personal idiom, and to the best of my knowledge, I'm the only one who merges symmetric cases using it. To minimize shock, here is the basic difference between the "classic" idiom and my own:

```

1 /* Classic binary search tree insertion */
2 struct cbt_node *cbt_insert ( struct cbt_node *root, int data )
3 {
4     if ( root == NULL )
5         root = make_node ( data );
6     else if ( data < root->data )
7         root->left = cbt_insert ( root->left, data );
8     else
9         root->right = cbt_insert ( root->right, data );
10
11     return root;
12 }
13
14 /* Julianne Walker's binary search tree insertion */
15 struct jsw_node *jsw_insert ( struct jsw_node *root, int data )
16 {
17     if ( root == NULL )
18         root = make_node ( data );
19     else {
20         int dir = root->data < data;
21         root->link[dir] = jsw_insert ( root->link[dir], data );
22     }
23
24     return root;
25 }

```

At this level there appears to be no savings, but if you add 50 lines of rebalancing code to both the left and right direction cases, it's easy to see that the classic insertion requires 100 extra lines, but my insertion only needs 50 because the two symmetric cases were merged into one. In the actual implementation this results in dir being the direction that we're going and !dir being the opposite direction. Once you get a feel for the idiom, it's not terribly hard to read code that uses it.

Back to the framework: to avoid constantly cluttering up the code

dary tests, a

impler. We

will assume that a leaf is black, because our leaves are null pointers and trying to fix a red violation with a null pointer scares the hell out of me. :-)

```

1 int is_red ( struct jsw_node *root )
2 {
3     return root != NULL && root->red == 1;
4 }

```

Finally, the rotations. Any tree balancing scheme will use rotations to change the structure without breaking any of the rules of binary search trees. However, in this case it's cleaner to also handle recoloring of the nodes that are rotated. A single rotation in a red black tree simply rotates the nodes as normal, then sets the old root to be red and

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

the new root to be black. A double rotation is just two single rotations. This is sufficient for insertion, but deletion is harder (what a shock). Despite that, we can still benefit from the color change during rotations in the deletion algorithm, so we will place it in the code for a single rotation, and the double rotation will use it implicitly:

```
1 struct jsw_node *jsw_single ( struct jsw_node *root, int dir )
2 {
3     struct jsw_node *save = root->link[dir];
4
5     root->link[dir] = save->link[dir];
6     save->link[dir] = root;
7
8     root->red = 1;
9     save->red = 0;
10
11     return save;
12 }
13
14 struct jsw_node *jsw_double ( struct jsw_node *root, int dir )
15 {
16     root->link[dir] = jsw_single ( root->link[dir], dir );
17     return jsw_single ( root, dir );
18 }
```

Don't get so anxious, we're not going to look at insertion just yet. Balanced trees are not trivial data structures, and red black trees are exceptionally tricky to get right because a small change in one part of the tree could cause a violation in a distant part of the tree. So it makes sense to have a little tester function to make sure that no violations have occurred. Why? Because it might seem like the algorithms work for small trees, but then they break on large trees and you don't know why. With a tester function, we can be confident that the algorithm works, provided we slam it with enough data into a big enough tree (too large of a case to test by hand). Since the tester would only be used for debugging, we can use recursion and expect it to be slow, which it is:

```
1 int jsw_rb_assert ( struct jsw_node *root )
2 {
3     int lh, rh;
4
5     if ( root == NULL )
6         return 1;
7     else {
8         struct jsw_node *ln = root->link[0];
9         struct jsw_node *rn = root->link[1];
10
11         /* Consecutive red links */
12         if ( is_red ( root ) ) {
13             if ( is_red ( ln ) || is_red ( rn ) ) {
14                 puts ( "Red violation" );
15             }
16         }
17     }
18 }
```

http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

125 captures

12 Feb 2007 - 8 Feb 2021

Go

SEP

OCT

JAN

29

2013

2021

2015

About this capture

```
19 int jsw_rb_assert ( int )
20 rh = jsw_rb_assert ( rn );
21
22 /* Invalid binary search tree */
23 if ( ( ln != NULL && ln->data >= root->data )
24     || ( rn != NULL && rn->data <= root->data ) )
25 {
26     puts ( "Binary tree violation" );
27     return 0;
28 }
29
30 /* Black height mismatch */
31 if ( lh != 0 && rh != 0 && lh != rh ) {
32     puts ( "Black violation" );
33     return 0;
34 }
35
36 /* Only count black links */
37 if ( lh != 0 && rh != 0 )
38     return is_red ( root ) ? lh : lh + 1;
39 else
40     return 0;
41 }
42 }
```

This algorithm is relatively simple. It walks over every node in the tree and performs certain tests on the node and its children. The first test is to see if a red node has red children. The second test makes sure that the tree is a valid binary search tree. The last test counts the black nodes along a path and ensures that all paths have the same black height. If `jsw_rb_assert` returns 0, the tree is an invalid red black tree. Otherwise it will return the black height of the entire tree. For convenience, a message will also print out telling you which violations occurred. :-)

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

Now we're ready to move on to insertion! So roll up your sleeves and get ready to mess about in the muck.

Bottom-up Insertion

When we want to insert into a red black tree, we immediately have a decision. Do we color a new node red or black? What kind of algorithm we write depends on this decision because of the violations that it could introduce. If the new node is black then inserting it into the tree always introduces a black violation. The rest of the algorithm would then need to concentrate on fixing the black violation without introducing a red violation. On the other hand, if the new node is red, there is a chance that it could introduce a red violation. The rest of the algorithm would then need to work toward fixing the red violation without introducing a black violation.

But wait! If the new node is red, and it's inserted as the child of a black node then no violations occur at all, whereas if the new node is black, a black violation always occurs. So the logical choice is to color the new node red because there is a possibility that insertion won't violate the rules at all. The same can't be said of inserting a black node. Red violations are also more localized and thus, easier to fix. Therefore, we will give new nodes the color red, and fix red violations during insertion. Laziness wins! How about a helper function that returns a new red node?

http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

Go

SEP

OCT

JAN

125 captures

12 Feb 2007 - 8 Feb 2021

29

2013

2021

2015

About this capture

```

6   rn->data = data;
7   rn->red = 1; /* 1 is red, 0 is black */
8   rn->link[0] = NULL;
9   rn->link[1] = NULL;
10  }
11
12  return rn;
13  }
```

Ah, finally. Insertion. Adding a node. Instead of just talking big about red black trees, we can play with the real thing now. Okay, let's insert a node into a basic binary search tree, that's the first order of business. Then we can walk back up and rebalance if there's a red violation. For simplicity (hah!) we'll use recursion to insert a new node and deny duplicates. Don't forget to notice the last bit of code to make the root black. That's important because it ensures the last red black tree rule and saves us from dealing with a red violation at the root. The code should be obvious, and if it isn't, you're in the wrong tutorial. Try Binary Search Trees I, third door to the left:

```

1  struct jsw_node *jsw_insert_r ( struct jsw_node *root, int data )
2  {
3      if ( root == NULL )
4          root = make_node ( data );
5      else if ( data != root->data ) {
6          int dir = root->data < data;
7
8          root->link[dir] = jsw_insert_r ( root->link[dir], data );
9
10         /* Hey, let's rebalance here! */
11     }
12
13     return root;
14 }
15
16 int jsw_insert ( struct jsw_tree *tree, int data )
17 {
18     tree->root = jsw_insert_r ( tree->root, data );
19     tree->root->red = 0;
20     return 1;
21 }
```

This is about as simple as it gets for a binary search tree, but since we know that we might need to rebalance back up the tree, naturally the code to rebalance would go after the recursive call. Remember those programming classes where you had to print numbers in reverse using recursion? The same principle applies here. We recurse down, down, down (down! sit! stay!), then insert a new node at the first leaf we get to (because the second leaf we get to would probably seg fault). Then the recursion rewinds on its own and we can take advantage of that.

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtree.aspx

Captured: Friday, March 12, 2021 15:53:10

So how do we test for a red violation? Simple, if any child on the path is red, test its children to see if one or the other is red. If so, fix the violation. Note that only one child will be red, otherwise something else is broken with the tree because a red violation would have to have already been present and the tree would already have been an invalid red black tree. But we'll assume that it's valid, because that's easier.

Okay, okay, but how do we FIX a red violation? That's equally simple! There are only three cases, and all of them are as easy as falling down. In the first case, we check both children of the node (remember that we don't do this for the new node). If both children are red, then the parent has to be black, so we can just flip the colors. The parent

http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtree.aspx

Go

SEP

OCT

JAN

29

125 captures


12 Feb 2007 - 8 Feb 2021

2013

2021


2015

About this capture




Okay, that takes care of the immediate problem, but if the parent's color changes then we risk a violation further up the tree. This is important, if the children of a node don't change color and the parent of a node doesn't change color, there's no way a red violation could propagate. So if 2 remained black in the above diagram, we could terminate. But since it changed from black to red, it's possible that its parent could also be red and there's another violation. So after this case we move up and look suspiciously around for problems.

The second case is if the node's (1 in this case) sibling (3 in this case) is not red. If so, the color flip trick won't work because that would increase the black height of the paths down to the right of 2. So we need to do more work, and that screams rotation. In the second case, 1's left child caused the violation, so we single rotate around 2 to the right, make 2 red, make 1 black, and call it a day:



This fixes the red violation, the subtree's root doesn't change color, so we can be sure that the red violation won't propagate upward. But it gets better! This rotation doesn't change the black height of either subtree, so the tree is now balanced and we're done with all rotations! But why does the black height not change? Notice that in the "before" diagram, the black height of 2's left subtree is 1 (including 2 itself) and the black height of 2's right subtree is 2. Now look at the "after" diagram. The left subtree's black height is still 1 and the right subtree's black height is still 2! Don't forget to ignore red nodes in the black height, they don't count.

Yes, good catch, that tree is not a valid red black tree. Why? Because the black heights are not the same. What we have here is a black violation! Durnit! But the good news is that this will never happen. The example was contrived to show you how the rotation works without having to draw a large tree, but that exact tree will never occur in the wild. Why? Because there was a black violation before 0 was inserted! In fact, the first rotation in a red black tree of ascending numbers would only be after the seventh number. Let's look at that case after 7 is inserted:



http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtree.aspx

Go

SEP

OCT

JAN

29

125 captures

12 Feb 2007 - 8 Feb 2021

2013

2021

2015


About this capture

Title: Eternally Confuzzled - Red Black Tree Tutorial


Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10


At this point a color flip will be made because 5, 4, and 6 meet the first case. A color flip could cause a red violation higher up, so we move up to 3 and see that it actually did cause a red violation at 3, and now the single rotation case applies. Let's look at the tree before the rotation is made:



The most important thing to notice right now is that there is NO black violation! Every path has two black nodes, so the only violation is the red violation at 3. When we rotate at 1, 3 becomes the new black root, and 1 becomes red. The resulting tree is as follows. Notice that there are no red violations, and all black heights are the same. So the single rotation case works:



It's pretty easy to see now why **jsw_single** sets colors the way it does. The final case is if instead of the left child being red, the right child is red. Well, if the right child is red then a single rotation won't cut it. This is where a double rotation is needed:



http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

125 captures

12 Feb 2007 - 8 Feb 2021

Go

SEP

NOV

JAN

29

2013

2014

2015

2016

2017

2018

2019

2020

2021

could introduce a black violation, so we continue the double rotation to avoid that case. See if you can find a valid red black tree where this case would exist and make sure that it actually works like we did above. In the meantime, we'll be here moving along without you. But don't worry, it's not like the tutorial will start to disappear if you stop along the way to work the exercises. :-)

The three cases have been looked at, now let's fill in that comment with the code to test for and fix all three cases. Remember that the new node is never rebalanced; the algorithm just returns the new node when it reaches a leaf and leaves rebalancing up to the nodes above. We make sure that if there's a violation, it's in the subtree. That way we can pull off the rotations without setting a status flag that reminds the algorithm to rotate further up. In fact, no status flag is needed at all! It's possible that we could do color flips all the way back up the tree even though only one single or double rotation might be made:

```

1 struct jsw_node *jsw_insert_r ( struct jsw_node *root, int data )
2 {
3     if ( !root == NULL )
  
```

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10


```

4  root = make_node ( "data" );
5  else if ( data != root->data ) {
6      int dir = root->data < data;
7
8      root->link[dir] = js_insert_r ( root->link[dir], data );
9
10     if ( is_red ( root->link[dir] ) ) {
11         if ( is_red ( root->link[dir] ) ) {
12             /* Case 1 */
13             root->red = 1;
14             root->link[0]->red = 0;
15             root->link[1]->red = 0;
16         }
17         else {
18             /* Cases 2 & 3 */
19             if ( is_red ( root->link[dir] > link[dir] ) )
20                 root = js_single ( root, !dir );
21             else if ( is_red ( root->link[dir] > link[!dir] ) )
22                 root = js_double ( root, !dir );
23         }
24     }
25 }
26
27 return root;
28 }
29
30 int js_insert ( struct js_tree *tree, int data )
31 {
32     tree->root = js_insert_r ( tree->root, data );
33     tree->root->red = 0;
34     return 1;
35 }

```

Okay, that was anticlimactic. Simple, elegant, that's red black insertion from the bottom-up! I recommend you go through an example that covers all of the cases, just to make sure that it works. Yes, we have the test function, but how do you know I'm not just messing with you? Well, you'll know after you grab about 30 random numbers and do your own execution trace!

Insertion really is the easiest part of red black trees. A red violation is sickly simple to test for and fix without causing problems elsewhere. The big problems come when you



Bottom-up Deletion

Deletion of a node from a red black tree is a pain in the ass. That's my official opinion too. I hate it and you'll learn to hate it. Especially the bottom-up algorithm. Especially the first way I'm going to describe it to you. >:-)

During insertion we had the option of selecting the color of a new node to make our lives easier. We forced a red violation and fixed it. Red violations are easy to fix, and we took full advantage of that to produce a truly elegant recursive algorithm. When you want to delete a node, you don't have the option of choosing its color. If it's red, that's great! Red nodes can be removed without any violations. Why? Well, I can't think of a way to introduce a red violation by removing a red node from a valid tree, but you're welcome to look for one. Since a red node doesn't participate in the black height, removing a red node has no effect on the black height. Therefore, removing a red node cannot violate the rules of a red black tree.

If the node is black, that's a problem. Removing a black node is sure to cause a black violation just like inserting a black node, which is why we avoided that during insertion, and it could very likely cause a red violation too! Ouch. Well, let's remove the node first, then think about how to fix violations. The following is a simple recursive deletion from a red black tree. If the node to be deleted has less than two children, we just replace it with its non-null child, or a null pointer if there are no children.

Conveniently enough, we might be able to get away without rebalancing at this point. If the node that we're going to delete is black and has one red child, we can color the child black and introduce no violations because we've reduced the case of removing a black node to the case of removing a red node, which is guaranteed to cause no violations. We'll call this case 0.

If the node has two children, we find its inorder predecessor and copy the predecessor's data into the node. Then we recursively delete the predecessor, since it's guaranteed to have at most one child. All in all, the algorithm is very pretty, even with the extra


Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsr_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

framework added to facilitate rebalancing (the simple recoloring after deletion, the status flag, and the mysterious `jsr_remove_balance`). Notice that the root is colored black at the end only if the tree is not empty:


```
1 struct jsr_node *jsr_remove_r ( struct jsr_node *root, int data, int *done )
2 {
3     if ( root == NULL )
4         *done = 1;
5     else {
6         int dir;
7
8         if ( root->data == data ) {
9             if ( root->link[0] == NULL || root->link[1] == NULL ) {
10                 struct jsr_node *save =
11                     root->link[root->link[0] == NULL];
12
13                 /* Case 0 */
14                 if ( is_red ( root ) )
15                     *done = 1;
16                 else if ( is_red ( save ) ) {
17                     save->red = 0;
18                     *done = 1;
19                 }
20             }
21             else {
22                 struct jsr_node *heir = root->link[0];
23                 while ( heir->link[1] != NULL )
24                     heir = heir->link[1];
25                 root->data = heir->data;
26                 data = heir->data;
27             }
28             dir = root->data < data;
29             root->link[dir] = jsr_remove_r ( root->link[dir], data, done );
30
31             if ( !*done )
32                 root = jsr_remove_balance ( root, dir, done );
33         }
34         return root;
35     }
36 }
37
38 int jsr_remove ( struct jsr_tree *tree, int data )
39 {
40     int done = 0;
41     tree->root = jsr_remove_r ( tree->root, data, &done );
42     if ( tree->root != NULL )
43         tree->root->red = 0;
44     return 1;
45 }
```



```
24     }
25     else {
26         struct jsr_node *heir = root->link[0];
27         while ( heir->link[1] != NULL )
28             heir = heir->link[1];
29         root->data = heir->data;
30         data = heir->data;
31     }
32     dir = root->data < data;
33     root->link[dir] = jsr_remove_r ( root->link[dir], data, done );
34
35     if ( !*done )
36         root = jsr_remove_balance ( root, dir, done );
37 }
38
39 int jsr_remove ( struct jsr_tree *tree, int data )
40 {
41     int done = 0;
42     tree->root = jsr_remove_r ( tree->root, data, &done );
43     if ( tree->root != NULL )
44         tree->root->red = 0;
45     return 1;
46 }
```

This time we really do need a status flag to tell the algorithm when to stop rebalancing. If `jsr_remove_balance` is called all of the way up the tree, it could get ugly. Since this function only removes external nodes (nodes without two children), we only have to worry about setting the flag after a removal (if the node is red or case 0 applies), and inside `jsr_remove_balance`.


It's obvious that the real work is done by `jsr_remove_balance`, but what does this helper function do? It handles all of the cases that could pop up after removing a black node, of course! And it's about as long as `jsr_remove_r`. :- (Let's look at the simple cases first. If we remove a node, and the node's sibling is black, and its children are both black, we have an easy case that propagates (* = the place we deleted):



In this case, we can simply recolor 1 (the sibling) to be red without introducing a red violation, but the possibility still exists that there's a black violation because we decreased the black height of 3's left subtree. This causes both of 3's subtrees to be the

same black height, but if 3 itself is a subtree, then its parent will see a black violation since the sibling subtree hasn't decreased in black height. However, if the parent node (3

125 captures



125 captures




Title: Eternally Confuzzled - Red Black Tree Tutorial


Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

sibling case looks at the inner child of the sibling. This child cannot be a null pointer (exercise: why?), so we simply test one of its children and act depending on which one is red. If the outer child is red, we perform a double rotation, color the new parent black, its right child black, and its left child red:



This restores the black height of the right subtree without changing the black height of the left subtree. When we have two red nodes to work with, it's easy to make one of them black after the rotation and leave the other to avoid violating the black height of the subtree that we took it from. The last case is if 3's right child is red. The good news is that we can reduce this to the previous case by a single rotation at 2, then a double rotation at 5 (the previous case) will give us the structure we want, with the same colors as above.



Before you start thinking of the terrible mess that `jsw_remove_balance` must be, let's look at it and see that it really isn't that bad. Yes, it's long, but none of the cases are overly difficult. The real savings come from noticing similarities between cases and setting colors outside of the cases so as to avoid repeating code, but we can do much better as you'll see shortly. Compare the cases described above with their translations into source code. Did we cover all of the cases?

http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtrees.aspx

Go

SEP

NOV

JAN

125 captures

12 Feb 2007 - 8 Feb 2021

29

2013

2014

2015

About this capture

```

6  if ( s != NULL && !is_red ( s ) ) {
7      /* Black sibling cases */
8      if ( !is_red ( s->link[0] ) && !is_red ( s->link[1] ) ) {
9          if ( is_red ( p ) )
10             *done = 1;
11             p->red = 0;
12             s->red = 1;
13         }
14         else {
15             int save = root->red;
16             if ( is_red ( s->link[!dir] ) )
17                 p = jsw_single ( p, dir );
18             else
19                 p = jsw_double ( p, dir );
20             p->red = save;
21             p->link[0]->red = 0;
22             p->link[1]->red = 0;
23             *done = 1;
24         }
25     }
26     else if ( s->link[dir] != NULL ) {
27         /* Red sibling cases */
28         struct jsw_node *r = s->link[dir];
29         if ( !is_red ( r->link[0] ) && !is_red ( r->link[1] ) ) {
30             p = jsw_single ( p, dir );
31             p->link[dir]->link[!dir]->red = 1;
32         }
33         else {

```

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx


Captured: Friday, March 12, 2021 15:53:10

```

37     if ( !is_red( r->link[dir]) )
38     {
39         s->link[dir] = jsw_single ( r, dir );
40         p = jsw_double ( p, dir );
41         s->link[dir]->red = 0;
42         p->link[dir]->red = 1;
43     }
44     p->red = 0;
45     p->link[dir]->red = 0;
46     *done = 1;
47 }
48
49 return p;
50 }

```

Traditionally, red black trees reduce the red sibling case to a black sibling case. Unless you're familiar with the red sibling case (which is why I showed it to you), the reasons behind this case reduction would be terribly confusing. But, because the sibling's parent and children have to be black, a single rotation will push the parent down on the side that we deleted, recolor it red, pull the sibling and its children up, and recolor the sibling black. This doesn't change the black height of the tree, it just reverses which side the violation is on. By pushing the entire violation down, we ensure that the new sibling (2 in this case) is black:



http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

Go

SEP

NOV

JAN

125 captures

12 Feb 2007 - 8 Feb 2021

tion

29

2013

2014

2015

2016

At this point

ABOUT THIS CAPTURE

need to move down along with the sibling. So * would still be the place we removed a node, but 2 would be the new sibling. The cases after this are identical, but this time we need to be careful to move down without losing anything so that we can move back up without any trouble:

```

1 struct jsw_node *jsw_remove_balance ( struct jsw_node *root, int dir, int *done )
2 {
3     struct jsw_node *p = root;
4     struct jsw_node *s = root->link[dir];
5
6     /* Case reduction, remove red sibling */
7     if ( is_red ( s ) ) {
8         root = jsw_single ( root, dir );
9         s = p->link[dir];
10    }
11
12    if ( s != NULL ) {
13        if ( !is_red ( s->link[0] ) && !is_red ( s->link[1] ) ) {
14            if ( is_red ( p ) )
15                *done = 1;
16            p->red = 0;
17            s->red = 1;
18        }
19        else {
20            int save = p->red;
21            int new_root = ( root == p );
22
23            if ( is_red ( s->link[dir] ) )
24                p = jsw_single ( p, dir );
25            else
26                p = jsw_double ( p, dir );
27
28            p->red = save;
29            p->link[0]->red = 0;
30            p->link[1]->red = 0;
31
32            if ( new_root )
33                root = p;
34            else
35                root->link[dir] = p;
36
37            *done = 1;
38        }
39    }
40
41    return root;
42 }

```

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

The subtle waltz of **root** and **p** could be confusing, but the only fix for that is an execution trace to make sure that you understand why **root** is used where it is and why **p** is used where it is. Remember that I said the case reduction was confusing (probably why nobody cares to explain it), but it simplifies the code greatly by removing half of the rebalancing work that we need to do.

Top-down Insertion

One of the primary benefits of red black trees is that both the Insertion and deletion algorithms can be written in one top-down pass, instead of an algorithm that walks down and then back up the tree.

http://www.eternallyconfuzzled.com/tuts/datastructures/js/tut_rbtrees.aspx

125 captures

12 Feb 2007 - 8 Feb 2021

Go

SEP

NOV

JAN

29

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

2025

at

29

SEP

NOV

JAN

29

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

12 Feb 2007 - 8 Feb 2021




This means that we don't need to use recursion any longer, or any other tricks that make the algorithms confusing. So let's look at a non-recursive top-down insertion.

One key concept that you need to take away from this tutorial is that the rules are the same, regardless of how you implement them. So there are still three cases for red black insertion: color flip, single rotation, and double rotation. Everything else is framework and details. However, for top-down insertion we don't have the benefit of riding the recursion wave. To recover from that deficiency we need to carefully save several pointers back up the tree. At the very least we have the current node along the path (which will be a null pointer at the end), its parent, and the parent's parent.


To avoid special cases, we'll also use a dummy tree root, where the actual root is the dummy's right link. This way we don't have to constantly test the special boundary case of rebalancing at the root. To make rotations simpler, we'll also add a pointer for the parent of the grandparent. So at the most we're saving the last three levels of the path.

For the color flip case, we test the current node's children. If they're red then the current node must be black, so we do a color flip. This could cause a red violation further up the tree (which is why we saved the parent of the current node). Now, because a color flip could immediately cause a violation, we go straight to the violation test between the parent and the current node. If both are red, we do a single or double rotation at the grandparent:

Color Flip



Single Rotation



Double Rotation



Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsv_tut_rbtrees.aspx

Captured: Friday, March 12, 2021 15:53:10

4,R ←-- p q,R 4,R

http://www.eternallyconfuzzled.com/tuts/datastructures/jsv_tut_rbtrees.aspx
Go
SEP NOV JAN
29
2013 2014 2015

125 captures
12 Feb 2007 - 8 Feb 2021
About this capture

1,B 3,B

Notice how a rotation is made around the grandparent, which is why we also save a pointer to the grandparent's parent. Otherwise it would be rather difficult to notify the tree of the changes we're making. The code to perform this magic is shorter than you would expect, but it doesn't follow conventional methods. All of the work is done in the walk down loop, including insertion, since inserting a new red node could cause a violation. So the loop breaks when the data is found either because a new node was inserted or because there was already a node with that data in the tree. As such, the following algorithm does not allow duplicates and also does not warn about a duplicate:

```

1 int jsv_insert ( struct jsv_tree *tree, int data )
2 {
3     if ( tree->root == NULL ) {
4         /* Empty tree case */
5         tree->root = make_node ( data );
6         if ( tree->root == NULL )
7             return 0;
8     }
9     else {
10        struct jsv_node head = {0}; /* false tree root */
11
12        struct jsv_node *g, *t;      /* Grandparent & parent */
13        struct jsv_node *p, *q;      /* Iterator & parent */
14        int dir = 0, last;
15
16        /* Set up helpers */
17        t = &head;
18        g = p = NULL;
19        q = t->link[1] = tree->root;
20
21        /* Search down the tree */
22        for ( ; ; ) {
23            if ( q == NULL ) {
24                /* Insert new node at the bottom */
25                p->link[dir] = q = make_node ( data );
26                if ( q == NULL )
27                    return 0;
28            }
29            else if ( is_red ( q->link[0] ) && is_red ( q->link[1] ) ) {
30                /* Color flip */
31                q->red = 1;
32                q->link[0]->red = 0;
33                q->link[1]->red = 0;
34            }
35
36            /* Fix red violation */
37            if ( is_red ( q ) && is_red ( p ) ) {
38                int dir2 = t->link[1] == q;
39
40                if ( q == p->link[last] )
41                    t->link[dir2] = jsv_single ( g, last );
42                else
43                    t->link[dir2] = jsv_double ( g, last );
44            }
45
46            /* Stop if found */
47            if ( q->data == data )
48                break;
49
50            last = dir;
51
52            q = p; p = q;
53            q = q->link[dir];
54        }
55
56        /* Update root */
57        tree->root = head.link[1];
58    }
59
60    /* Make root black */
61    tree->root->red = 0;
62
63    return 1;
64 }

```

http://www.eternallyconfuzzled.com/tuts/datastructures/jsv_tut_rbtrees.aspx
Go
SEP NOV JAN
29
2013 2014 2015

125 captures
12 Feb 2007 - 8 Feb 2021
About this capture

```

56     q = p; p = q;
57     q = q->link[dir];
58 }
59
60 /* Update root */
61 tree->root = head.link[1];
62 }
63
64 /* Make root black */
65 tree->root->red = 0;
66
67 return 1;
68 }

```


Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx


Captured: Friday, March 12, 2021 15:53:10

Granted, this top-down insertion isn't as pretty as the recursive bottom-up insertion, but it takes full advantage of the properties of red black trees, and avoids the potential problems of recursion. It's also theoretically more efficient. :-) The good news is that deletion is easier top-down than bottom-up, and we can use the lessons gained from insertion to make implementing it easier. Let's take a look.


Top-down Deletion

Bottom-up deletion was a pain because we needed to recover from a black violation, which is the harder of the two violations to fix. However, the deletion was trivial if we removed a red node because deleting a red node can't violate any of the rules of a red black tree. If we could somehow guarantee that the node to be deleted was red, it would simplify deletion greatly. The good news is that we can do just that! By forcing a red node into the tree at the top and pushing it down using color flips and rotations, we can ensure that a red node will be deleted, always.

There are only four cases for pushing a red node down the tree without a black violation. The first case is a simple reverse color flip. If a node and its sibling are black, and all four of their children are black, make the parent black and both children red:



This is basically a reverse color flip from what we did during insertion. It pushes the red nodes down rather than forcing them up, all without changing the black height of the tree. Cool. :-) The next case is the dreaded red sibling case, which we will quickly reduce to nothing with a single rotation, using the tricks we learned from bottom-up deletion. This case is very intuitive now that we don't have to change directions to get it to work. Notice how the black heights don't change, nor does the color of the parent. Since we're moving down already, pushing the sibling down is just what the doctor ordered:



http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx

125 captures

12 Feb 2007 - 8 Feb 2021

Go

SEP

NOV

JAN

29


2013

2014


2015

ABOUT US SECURITY

The final two cases are, you guessed it, single and double rotations depending on the color of the sibling's children. The cases are (should be!) familiar by now, so we'll just go over them briefly before looking at the code. If the right child is red, we perform a double rotation. However, in this case the color changes of `jsw_single` won't quite cut it. To be thorough, we'll force the correct coloring for all affected nodes:



As always, we have to make sure that the color of the parent doesn't change and the black heights of both subtrees remain the same as they were originally. These two guidelines avoid both red and black violations. If the left child of the sibling is red then a single rotation is more than enough to do the same thing. Since the final colors are the same, so the recoloring of this case can be lumped together with the previous case:




Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtrees.aspx


Captured: Friday, March 12, 2021 15:53:10

Now for the fun part. This algorithm will use a sneaky deletion by copying technique that saves the node to be deleted, and just keeps on going through the loop until the inorder predecessor is found. At that point the walk down loop terminates, the data items are copied and the external node is spliced out of the tree. This is as opposed to a naive deletion by copying that breaks the search off when the node is found, then has an extra step that explicitly looks for the predecessor. On the way down, of course, we look for and handle the cases shown above:

```
1 int jsw_remove ( struct jsw_tree *tree, int data )
2 {
3     if ( tree->root != NULL ) {
4         struct jsw_node head = {0}; /* False tree root */
5         struct jsw_node *q, *p, *g; /* Helpers */
6         struct jsw_node *f = NULL; /* Found item */
7         int dir = 1;
8
9         /* Set up helpers */
10        q = &head;
11        p = NULL;
12        q->link[1] = tree->root;
13
14        /* Search and push a red down */
15        while ( q->link[dir] != NULL ) {
16            int last = dir;
17
18            /* Update helpers */
19            g = p, p = q;
20            q = q->link[dir];
```



```
26        /* Push the red node down */
27        if ( !is_red ( q ) && !is_red ( q->link[dir] ) ) {
28            if ( !is_red ( q->link[!dir] ) )
29                p = p->link[last] = jsw_single ( q, dir );
30            else if ( !is_red ( q->link[!dir] ) ) {
31                struct jsw_node *s = p->link[last];
32
33                if ( s != NULL ) {
34                    if ( !is_red ( s->link[!last] ) && !is_red ( s->link[last] ) ) {
35                        /* Color flip */
36                        p->red = 0;
37                        s->red = 1;
38                        q->red = 1;
39                    }
40                    else {
41                        int dir2 = q->link[1] == p;
42
43                        if ( !is_red ( s->link[last] ) )
44                            q->link[dir2] = jsw_double ( p, last );
45                        else if ( !is_red ( s->link[!last] ) )
46                            q->link[dir2] = jsw_single ( p, last );
47
48                        /* Ensure correct coloring */
49                        q->red = q->link[dir2]->red = 1;
50                        q->link[dir2]->link[0]->red = 0;
51                        q->link[dir2]->link[1]->red = 0;
52                    }
53                }
54            }
55        }
56    }
57 }
58
59 /* Replace and remove if found */
60 if ( f != NULL ) {
61     f->data = q->data;
62     p->link[p->link[1] == q] =
63     q->link[q->link[0] == NULL];
64     free ( q );
65 }
66 }
```



```
72     return 1;
73 }
74 }
```

It's short and sweet, though the extreme level of indentation is a little distasteful for some. As an exercise, try to shrink that down. :-) As with insertion we used a dummy root to avoid special cases, and a few helper variables to point up the tree. The majority of the complexity comes from our rebalancing cases, and if you remove them, you'll

Title: Eternally Confuzzled - Red Black Tree Tutorial

Link: https://web.archive.org/web/20141129024312/http://www.eternallyconfuzzled.com/tuts/datastructures/jsw_tut_rbtree.aspx

Captured: Friday, March 12, 2021 15:53:10

discover a very clean deletion algorithm for a basic binary search tree. Be sure to trace the execution of this function to make sure that you understand how it works, because it's a very nifty algorithm. Though I'm a little biased because I've never seen it used elsewhere, which means I may be the first to discover it. :-)

That's top-down deletion, and it can be compared to bottom-up deletion in the same ways as insertion, but now we can also consider that the top-down code is shorter and (I think) simpler. Red black trees are generally cleaner and easier to understand if written in a top-down manner, but unfortunately, due to the fact that the only well known resource that gives code for deletion is CLR, the most common implementations are still bottom-up. Oh well.

Conclusion

Red black trees are interesting beasts. They're believed to be simpler than AVL trees (their direct competitor), and at first glance this seems to be the case because insertion is a breeze. However, when one begins to play with the deletion algorithm, red black trees become very tricky. However, the counterweight to this added complexity is that both insertion and deletion can be implemented using a single pass, top-down algorithm. Such is not the case with AVL trees, where only the insertion algorithm can be written top-down. Deletion from an AVL tree requires a bottom-up algorithm.

So when do you use a red black tree? That's really your decision, but I've found that red black trees are best suited to largely random data that has occasional degenerate runs, and searches have no locality of reference. This takes full advantage of the minimal work that red black trees perform to maintain balance compared to AVL trees and still allows for speedy searches.

Red black trees are popular, as most data structures with a whimsical name. For example, in Java and C++, the library map structures are typically implemented with a red black tree. Red black trees are also comparable in speed to AVL trees. While the balance is not quite as good, the work it takes to maintain balance is usually better in a red black tree. There are a few misconceptions floating around, but for the most part the hype about red black trees is accurate.

This tutorial described the red black abstraction for balancing binary search trees. Both insertion and deletion were covered in their top-down and bottom-up forms, and a full implementation for each form was given. For further information on red black trees, feel free to surf over to Ben Pfaff's excellent online book about libavl (I wish I had it when figuring this stuff out!), or pick up "Introduction to Algorithms" by Cormen, Leiserson, and Rivest.

From the twisted mind of Julianne Walker

[Top](#)